

REMARKS

Reconsideration of the above-identified patent application in view of the amendments above and the remarks following is respectfully requested.

Claims 1 and 3-80 are in this case. Claims 41-80 were withdrawn by the Examiner from consideration under traverse of Applicant as drawn to a non-elected invention. Claims 1 and 3-40 have been rejected under § 103(a). Dependent claims 6, 11 and 13 have been canceled. Independent claims 1 and 28 and dependent claims 5, 7, 14, 15, 21 and 22 have been amended.

The claims before the Examiner are directed toward an interface adapter for a packet network.

In one embodiment of the present invention, the interface adapter includes a memory interface, a plurality of execution engines, a scheduling processor, a plurality of gather engines and switching circuitry. The memory interface is coupled to a memory. The execution engines read, from the memory via the memory interface, work items that have been assigned to them by the scheduling processor, and generate corresponding gather entries. Each work item belongs to a respective one of multiple transport service instances. The scheduling processor receives the transport service instances from a host of the interface adapter and enters each transport service instance in a respective one of a multiplicity of scheduling queues according to the transport service instance's respective class of service. The scheduling processor also assigns the work items to the execution engines by selecting the corresponding transport service instances from the scheduling queues for service in accordance with the transport service instances' classes of service, so that the execution engines read the work items from the memory and the gather engines receive the gather entries from the execution engines and generate corresponding packets. The switching

circuitry couples the execution engines to the gather engines so as to submit the gather entries to the gather engines.

In another embodiment of the present invention, the interface adapter includes a memory interface, a plurality of execution engines, a scheduling processor and a send data unit. The memory interface is coupled to a memory in which are stored work items belonging to transport service instances that are assigned to respective classes of service. The execution engines read the work items from the memory via the memory interface and generate corresponding gather entries. The scheduling processor maintains scheduling queues for the transport service instances according to the respective classes of service of the transport service instances, receives the transport service instances from a host of the interface adapter, and selects transport service instances from the heads of the queues for servicing by the execution engines. The send data unit generates packets responsive to the gather entries.

§ 103(a) Rejections – Tzeng et al. ‘604 in view of Gasbarro et al. ‘004

The Examiner has rejected claims 1, 3-8, 10-14, 17, 22, 24-26, 28-31, 34 and 37 under § 103(a) as being unpatentable over Tzeng et al., US Patent No. 6,912,604 (henceforth, “Tzeng et al. ‘604”) in view of Gasbarro et al., US Patent No. 6,948,004 (henceforth, “Gasbarro et al. ‘004”). The Examiner’s rejection is respectfully traversed.

Tzeng et al. ‘604 teach a host channel adapter (HCA) **12**, for sending packets on an InfiniBand™ network **10**, in which some link layer operations are performed before the transport layer operations and the rest of the link layer operations are performed after the transport layer operations. Specifically, host channel adapter **12** includes a pre-link module **40**, a transport service module **42** and a post-link module **44**. Pre-link module **40** receives work queue elements (WQEs) from a host and stores

the WQEs in a WQE FIFO 50. A pre-link process module 54 assigns WQEs from WQE FIFO 50 to virtual lane FIFOs 52 according to the WQEs' service levels as determined from service level information in the WQEs themselves. A VL arbitration module 60 sends WQEs from virtual lane FIFOs 52 to transport service module 42 according to the respective priorities of virtual lane FIFOs 52. Transport service module 42 matches up the WQEs with the corresponding queue pairs (QPs) and sends the WQEs to post-link module 44 that generates the corresponding packets.

The Examiner has cited Gasbarro et al. '004 primarily for the supposed teaching therein of the coupling of execution engines to a memory interface to read work items. In addition, the Examiner's Response to Arguments on pages 2-7 of the present Office Action are based exclusively on the Examiner's interpretation of Tzeng et al. '604 and not on the Examiner's interpretation of Gasbarro et al. '004. Therefore, the following remarks are directed primarily at Tzeng et al. '604.

Briefly, as noted in the responses filed July 16, 2006 and October 27, 2006 to the Office Action mailed January 17, 2006, and as noted in the response filed January 6, 2009 to the Office Action mailed July 11, 2008, the difference between the teachings of Tzeng et al. '604 and the present invention is that Tzeng et al. '604 perform service level classification of WQEs, whereas the present invention performs service level classification of QPs (or, more generally, of transport service instances). HCA 12 of Tzeng et al. '604 receives WQEs from the host, classifies the WQEs according to service levels as determined from service level information in the WQEs themselves, matches up WQEs with QPs only after the WQEs have been classified according to service levels, and generates corresponding packets. The present invention uses a scheduling processor to receive transport service instances from the host and to assign the transport service instances to schedule queues according to

service level before ever reading the associated WQEs. The present invention then matches up the transport service instances with their WQEs (or, more generally, with their work items) in execution engines, and generates the corresponding packets in gather engines.

In the Response to Arguments, on page 4 of the present Office Action, the Examiner argues that

...the argued feature of service classification of WQEs versus service classification of QPs is...nowhere to be found in the claim recitation.

Applicant respectfully disagrees. Claim 28, even before the present amendments, recites

a scheduling processor, adapted to maintain a multiplicity of scheduling queues corresponding to the classes of service...and further adapted to enter the transport service instances...in the scheduling queues according to the classes of service of the instances... (emphasis added)

This claim language clearly recites service level classification of transport service instances (QPs), not service level classification of work items (WQEs).

Then, on page 6 of the present Office Action, the Examiner cites Tzeng et al. '604 column 4 lines 28-46 as teaching the limitation, recited in claim 1 prior to the present amendments, of

wherein each of the work items belongs to a respective transport service instance among multiple transport service instances served by the adapter, and wherein the scheduling processor is adapted to assign the work items to the execution engines by selecting the respective transport service instances for service

The cited portion of Tzeng et al. '604 teaches no such thing. In particular, column 4 line 39, cited by the Examiner, states explicitly that

The WQE includes service level (SL) information...

and then column 4 lines 55-58, not cited by the Examiner, states explicitly that

In particular, the pre-link process module 54 retrieves a WQE from the WQE FIFO 50, and determines the corresponding virtual lane based on the service layer specified within the WQE. (emphasis added)

rather than based on a service level specified by the QP to which the WQE belongs.

Later, in column 6 lines 16-19, Tzeng et al. '604 teach

The queue pair attributes management module 66 also includes service submodules 68, each configured for managing a corresponding transport service type based on a corresponding received WQE from the pre-link module 40. (emphasis added)

It thus is abundantly clear that Tzeng et al. '604 schedule WQEs for processing directly according to the service levels specified in the WQEs themselves, and not indirectly according to service levels (classes of service) specified by their QPs as in the present invention.

To more clearly distinguish the present invention from the prior art cited by the Examiner, independent claims 1 and 28 have been amended as follows.

In claim 1, the limitation that each work item belongs to a respective transport service instance has been moved to the recitation of the execution engines, for clarity.

In claim 1, the limitation that the scheduling processor assigns the work items to the execution engine by selecting the respective transport service instance for service has been moved to the recitation of the scheduling processor, for clarity.

The limitation, that the scheduling processor enters each transport service instances in a respective one of a multiplicity of scheduling queues according to a respective class of service of the transport service instance, has been added to the recitation of the scheduling processor in claim 1. This limitation corresponds to the limitations, recited in claim 28, that the scheduling processor is adapted to maintain a multiplicity of scheduling queues corresponding to the classes of service and that the scheduling processor is adapted to enter the transport service instances in the scheduling queues according to the classes of service of the instances.

The limitation, that the selecting, by the scheduling processor, of the transport service instances is from the scheduling queues in accordance with the transport service instances' respective classes of service, has been added to the recitation of the scheduling processor in claim 1. This limitation corresponds to the limitations, recited in claim 28, that the scheduling processor is adapted to maintain a multiplicity of scheduling queues corresponding to the classes of service and that the scheduling processor selects the instances from the heads of the queues for assignment to the execution engines.

The limitation, that the scheduling processor is adapted to receive the transport service instances from a host of the interface adapter, has been added to both claim 1 and claim 28. Support for this limitation is found in the specification as filed at least on page 19 lines 17-20:

When the host processor rings its assigned doorbell 50, a doorbell handler 54 enters the corresponding QP in one of a plurality of schedule queues 56 that it maintains.

The limitation, that the scheduling processor assigns the work items to the execution engines for reading the work items from the memory, has been added to both claim 1 and claim 28. Support for this limitation is found in the specification as filed at least on page 20 lines 26-27:

The assigned execution engine fetches the WQEs of the selected QP from memory 38.

The claims that depend from claim 1 have been amended for consistency with claim 1 as now amended. For example, claim 6, that recites limitations now recited in claim 1, has been canceled, and claim 7 has been amended to depend directly from claim 1.

In order for claims 1 and 28 as now amended to be unpatentable over the Examiner's proposed combination of Tzeng et al. '604 and Gasbarro et al. '004, these

references must teach or suggest every recited limitation. As the Board of Patent Appeal and Interferences has recently confirmed in *In re Wada and Murphy*, Appeal 2007-3733,

When determining whether a claim is obvious, an examiner must make “a searching comparison of the claimed invention – *including all its limitations* – within the teaching of the prior art”. *In re Orchiai*, 71 F.3d 1565, 1572 (Fed. Cir. 1995) (emphasis added). Thus, “Obviousness requires a suggestion of all limitations in a claim.” *CFMT, Inc. v. Yieldup Intern. Corp.*, 349 F.3d 1333, 1342 (Fed. Cir. 2003) (citing *In re Royka*, 490 F.2d 981, 985 (CCPA 1974)).

Neither Tzeng et al. ‘604 nor Gasbarro et al. ‘004 teach, hint or suggest a scheduling processor, of an interface adapter, that receives QPs *as such* from the host of the interface adapter, and that assigns the QPs (rather than the associated WQEs) to execution engines, for reading of the associated WQEs from the host’s memory by the execution engines followed by generation of gather entries by the execution engines, according to the classes of service of the QPs. In particular, in Tzeng et al. ‘604, transport service module 42, whose submodules 68 the Examiner has identified with the execution engines recited in claims 1 and 28, does not read the WQEs from system memory 26, but instead receives the WQEs from pre-link module 40. This feature of Tzeng et al. ‘604 is stated explicitly in the above citation from column 6 lines 16-19 and also in column 6 lines 41-43:

Hence, the transport service module 42, upon receiving a WQE from the pre-link module 40, supplies the WQE to the appropriate submodule 68 for processing... (emphasis added)

With independent claims 1 and 28 allowable in their present form it follows that claims 3-8, 10-14, 17, 22, 24-26, 29-31, 34 and 37 that depend therefrom also are allowable.

**§ 103(a) Rejections – Tzeng et al. ‘604 in view of Gasbarro et al. ‘004 and further
in view of Pettey et al. ‘712**

The Examiner has rejected claims 9, 27 and 32 under § 103(a) as being unpatentable over Tzeng et al. ‘604 in view of “Gasbarro et al., US Patent No. 6,594,712” and further in view of Pettey et al., US Patent No. 6,594,712 (henceforth, “Pettey et al. ‘712”). Applicant presumes that the Examiner intended to reject claims 9, 27 and 32 as being unpatentable over Tzeng et al. ‘604 in view of Gasbarro et al. ‘004 and further in view of Pettey et al., US Patent No. 6,594,712 (henceforth, “Pettey et al. ‘712”). The Examiner’s rejection is respectfully traversed.

It is demonstrated above that independent claims 1 and 28 are allowable in their present form. It follows that claims 9, 27 and 32 that depend therefrom also are allowable.

§ 103(a) Rejections – Tzeng et al. ‘604 in view of AAPA

The Examiner has rejected claims 18 and 19 under § 103(a) as being unpatentable over Tzeng et al. ‘604 in view of Applicant’s Admitted Prior Art. The Examiner’s rejection is respectfully traversed.

It is demonstrated above that independent claim 1 is allowable in its present form. It follows that claims 18 and 19 that depend therefrom also are allowable.

**§ 103(a) Rejections – Tzeng et al. ‘604 in view of Gasbarro et al. ‘004 and further
in view of Parthasarathy et al. ‘916**

The Examiner has rejected claims 15, 20, 21, 23, 35, 36 and 38 under § 103(a) as being unpatentable over Tzeng et al. ‘604 in view “Gasbarro et al., US Patent No. 6,594,712” and further in view of Parthasarathy et al., US Patent No. 6,831,916 (henceforth, “Parthasarathy et al. ‘916”). Applicant presumes that the Examiner

intended to reject claims 15, 20, 21, 23, 35, 36 and 38 as being unpatentable over Tzeng et al. '604 in view of Gasbarro et al. '004 and further in view of Parthasarathy et al. '916. The Examiner's rejection is respectfully traversed.

It is demonstrated above that independent claims 1 and 28 are allowable in their present form. It follows that claims 15, 20, 23, 35, 36 and 38 that depend therefrom also are allowable.

§ 103(a) Rejections – Tzeng et al. '604 in view of Gasbarro et al. '004, in view of Pettley et al. '712 and further in view of Parthasarathy et al. '916

The Examiner has rejected claim 33 under § 103(a) as being unpatentable over Tzeng et al. '604 in view of Gasbarro et al. '004, in view of Pettley et al. '712 and further in view of Parthasarathy et al. '916. The Examiner's rejection is respectfully traversed.

It is demonstrated above that independent claim 28 is allowable in its present form. It follows that claim 33 that depends therefrom also is allowable.

Applicant notes in passing that in rejecting claim 33, the Examiner has misidentified the transport service instance recited in claim 33 with WQEs, on lines 4 and 19 of page 32 of the present Office Action. The transport service instances recited in the claims of the above-identified patent application are generalizations of InfiniBand QPs, not of InfiniBand WQEs. The identification of QPs as special cases of transport service instances can be found in the specification on page 2 lines 5-9:

Client processes running on the host communicate with the transport layer of the IB fabric by manipulating a transport service instance, known as a "queue pair" (QP), made up of a send work queue and a receive work queue.

**§ 103(a) Rejections – Tzeng et al. ‘604 in view of Gasbarro et al. ‘004 and further
in view of Grun ‘591**

The Examiner has rejected claim 16 under § 103(a) as being unpatentable over Tzeng et al. ‘604 in view of “Gasbarro et al., US Patent No. 6,594,712” and further Grun, US Patent No. 6,272,591. Applicant presumes that the Examiner intended to reject claim 16 as being unpatentable over Tzeng et al. ‘604 in view of Gasbarro et al. ‘004 and further in view of Grun, US Patent No. 6,272,591. The Examiner’s rejection is respectfully traversed.

It is demonstrated above that independent claim 1 is allowable in its present form. It follows that claim 16 that depends therefrom also is allowable.

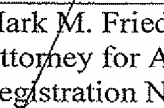
**§ 103(a) Rejections – Tzeng et al. ‘604 in view of Gasbarro et al. ‘004 and further
in view of Snyder II et al. ‘830**

The Examiner has rejected claims 39 and 40 under § 103(a) as being unpatentable over Tzeng et al. ‘604 in view of “Gasbarro et al., US Patent No. 6,594,712” and further in view of Snyder II et al., US Patent No. 6,888,830. The Examiner’s rejection is respectfully traversed. Applicant presumes that the Examiner intended to reject claims 39 and 40 as being unpatentable over Tzeng et al. ‘604 in view of Gasbarro et al. ‘004 and further in view of Snyder II et al., US Patent No. 6,888,830. The Examiner’s rejection is respectfully traversed.

It is demonstrated above that independent claim 28 is allowable in its present form. It follows that claims 39 and 40 that depend therefrom also are allowable.

In view of the above amendments and remarks it is respectfully submitted that independent claims 1 and 28, and hence dependent claims 3-27 and 29-40 are in condition for allowance. Prompt notice of allowance is respectfully and earnestly solicited.

Respectfully submitted,



Mark M. Friedman
Attorney for Applicant
Registration No. 33,883

Date: November 11, 2009